

A Discontinuous Neural Network for Non-Negative Sparse Approximation

Martijn Arts, Marius Cordts, Monika Gorin, Marc Spehr, and Rudolf Mathar

Abstract—This paper investigates a discontinuous neural network which is used as a model of the mammalian olfactory system and can more generally be applied to solve non-negative sparse approximation problems. By inherently limiting the systems integrators to having non-negative outputs, the system function becomes discontinuous since the integrators switch between being inactive and being active. It is shown that the presented network converges to equilibrium points which are solutions to general non-negative least squares optimization problems. We specify a Caratheodory solution and prove that the network is stable, provided that the system matrix has full column-rank. Under a mild condition on the equilibrium point, we show that the network converges to its equilibrium within a finite number of switches. Two applications of the neural network are shown. Firstly, we apply the network as a model of the olfactory system and show that in principle it may be capable of performing complex sparse signal recovery tasks. Secondly, we generalize the application to include non-negative sparse approximation problems and compare the recovery performance to a classical non-negative basis pursuit denoising algorithm. We conclude that the recovery performance differs only marginally from the classical algorithm, while the neural network has the advantage that no performance critical regularization parameter has to be chosen prior to recovery.

Index Terms—Discontinuity, optimization, NNLS, non-negative sparse approximation, olfactory system.

I. INTRODUCTION

SPARSITY is a concept that has recently generated a considerable amount of interest in mathematics, computer science and electrical engineering. In essence, it states that many signals may be represented by only a small number of non-zero coefficients using a suitable basis or dictionary. The field of compressed sensing exploits knowledge about the sparsity which allows to recover signals from a smaller number of measurements compared to classical methods [1], [2]. This is done by solving optimization problems that penalize non-sparse solution candidates.

In this work, we study a sparse signal recovery problem faced by the mammalian main olfactory system, namely the decomposition of a mixture of chemical stimuli from noisy measurements. The result of our modeling is a discontinuous

neural network that solves a quadratic program (QP) in order to recover the signal of interest from the measurements. We later show that this network can be applied more generally to solve non-negative sparse approximation problems, which do not have to be related to the biological application.

The connection between neural networks and optimization problems has been investigated for several decades now. A particularly interesting class, called (continuous) *Hopfield Neural Networks* (HNNs), was proposed by Hopfield in the 1980s [3], [4]. These networks consider either weighted excitatory or inhibitory feedback between several “neurons”. The neurons themselves are modeled as amplifiers with a potentially arbitrary input-output function. A sufficient condition for stability of such a network is that the synaptic weights are symmetric and feedback loops, i.e., the output of a neuron is connected to its own input, are not present. Such networks can be used to solve a variety of optimization problems, see [5] for a review. In the context of QPs, known networks used to solve (bounded) QPs such as [6], [7] are functionally related to the network presented in this work. These networks, however, differ fundamentally in structure and use other methods to investigate the convergence behavior compared to our approach. Specifically for sparse approximation scenarios, Locally Competitive Algorithms (LCAs) [8] were proposed. They solve optimization problems with a quadratic objective function, regularized by a cost function that is applied to the output of the integrators. Depending on the choice of the cost function, these neural networks solve the well known basis pursuit denoising problem (see also (3)) or an approximation of the original non-convex compressed sensing problem, where each non-zero entry is penalized independent of its actual amplitude (cf. (1)). The convergence behavior of the LCAs was studied in detail in [9].

The contributions of this paper are summarized in the following. We study a neural network in the form of a dynamic system where the integrators of the system are limited to having non-negative outputs. This results in discontinuities of the system function, since the integrators switch between being active and inactive depending on the value of their output. First, we show that the neural networks equilibrium points are the minimizers of a general non-negative least squares optimization problem, then we specify a Caratheodory solution for the systems discontinuous differential equations. Having defined the Caratheodory solution, we use Lyapunov’s direct method to show that the neural network is asymptotically stable, provided that the system matrix has full column-rank. Under a mild condition on the equilibrium point, we show that the network converges to this point within a finite number of switches. Subsequently, we discuss that the non-negative integrators can be generalized to arbitrary lower and

M. Arts and R. Mathar are with the Institute for Theoretical Information Technology, RWTH Aachen University, Aachen, Germany.

M. Cordts is now with the Environment Perception Dept. at Daimler R&D, Sindelfingen, Germany.

M. Gorin and M. Spehr are with the Department of Chemosensation, Institute for Biology II, RWTH Aachen University, Aachen, Germany.

This work was partly supported by the Deutsche Forschungsgemeinschaft (DFG) SPP 1395: Informations- und Kommunikationstheorie in der Molekularbiologie (InKoMBio), Grants: MA 1184/20-2 & SP 724/8-1.

Please direct your e-mail correspondence to: arts@ti.rwth-aachen.de

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

upper limits, so that the equilibrium point corresponds to the solution of a bounded least squares problem. Two different applications of the neural network are presented. The first is the original application, which served as the foundation for the presented neural network. It is derived as a model of the mammalian olfactory system and we investigate the performance of a stimulus mixture decomposition scenario. Based on the network performance, we conclude that the olfactory system may indeed recover the composition of chemical stimuli from a noisy mixture of neural responses. The second application deals with the more general class of non-negative sparse approximation problems. Here, we compare the recovery performance of the neural network with a non-negative basis pursuit denoising method and show that for this type of application the stability of the network remains unaffected even if the system matrix does not have full column-rank. From this investigation, we conclude that the neural network performance differs only marginally compared to the “classical” recovery method. However, the network has the additional advantage that no performance critical regularization parameter has to be determined prior to recovery.

An overview of the structure of this paper is given in the following. In Section II-A, we give the necessary background about compressed sensing / sparse approximation and summarize relevant results on the when dealing with non-negative sparse recovery problems. A concise overview of the stability theory of (discontinuous) dynamic systems can be found in Section II-B. The discontinuous neural network which is considered in this work is presented in Section III and relevant notation is defined. A formal investigation of the convergence properties and stability analysis of the neural network is given in Section IV. We present two applications of the presented neural network. Firstly, it is used as a model of the mammalian olfactory system, where it is shown that such a network would be able to decompose a mixture of chemical stimuli (Section V). Secondly, in Section VI, we show that the network is more generally able to solve non-negative sparse approximation problems. Section VII concludes the paper.

II. BACKGROUND

A. Compressed Sensing / Sparse Approximation

Compressed sensing is a relatively new research field that deals with recovering signals from a lower number of linear measurements compared to classical methods of signal recovery [1], [2]. The field relies on a concept called *sparsity*, i.e., the assumption that the signal can be represented by only a few nonzero coefficients in a suitable basis. A signal $\mathbf{x} \in \mathbb{R}^N$ is called s -sparse if at most s coefficients of the vector \mathbf{x} are nonzero, that is, if $\|\mathbf{x}\|_0 \leq s$. Here, $\|\mathbf{x}\|_0 := \#\{j : x_j \neq 0\}$ denotes the ℓ_0 pseudo-norm. Note that this “norm” does not fulfill the homogeneity property and therefore does not represent a proper norm.

The task in compressed sensing is to recover an s -sparse signal \mathbf{x} from a measured signal $\mathbf{b} \in \mathbb{R}^M$, obtained by a linear measurement process described by $\mathbf{b} = \mathbf{A}\mathbf{x}$ with $\mathbf{A} \in \mathbb{R}^{M \times N}$. Here, the number of measurements is small compared to the dimension of the original signal, i.e., $M \ll N$. Ideally,

one would like to recover \mathbf{x} using the following optimization problem:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \|\mathbf{x}\|_0 \\ & \text{subject to} && \mathbf{b} = \mathbf{A}\mathbf{x}. \end{aligned} \quad (1)$$

However, solving problem (1) is NP-hard and thus intractable [2]. If the matrix \mathbf{A} fulfills certain conditions however, e.g., the so-called null space property or the restricted isometry property, the minimizer of the ℓ_1 -minimization problem (also called basis pursuit)

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \|\mathbf{x}\|_1 \\ & \text{subject to} && \mathbf{b} = \mathbf{A}\mathbf{x}, \end{aligned} \quad (2)$$

is also the minimizer of problem (1) and hence may recover the desired signal vector exactly [1], [2]. In problem (2), the ℓ_1 -norm, which is the convex relaxation of the ℓ_0 -“norm”, was used as the objective. Note, that we confined \mathbf{x} , \mathbf{A} and \mathbf{b} to the real numbers but the results summarized above also hold in the complex setting.

In practical scenarios, virtually all measurements are perturbed by some form of noise. Thus, in practice the measurement process may be modeled as $\mathbf{b} = \mathbf{A}\mathbf{x} + \varepsilon$ where $\varepsilon \in \mathbb{R}^M$ denotes the noise component. In this case, taking $M = N$ measurements will not guarantee exact recovery of the desired signal vector anymore. Hence, one would typically take more measurements ($M > N$) and try to recover a good approximation of the signal vector through least-squares minimization. Likewise, in compressed sensing, one cannot hope to recover the exact signal vector when measurements are noisy. Consequently, the linear constraint $\mathbf{b} = \mathbf{A}\mathbf{x}$ in problem (2) is too strict since it does not allow for deviations caused by noise. To cope with this, one may replace the constraint in problem (2) by a quadratic constraint $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 \leq c_\varepsilon$ which allows for deviations within a sphere with constant radius. To choose c_ε meaningfully, it is necessary to have prior knowledge about the noise power. Another method is regularizing the objective function, as done in basis pursuit denoising (BPDN):

$$\underset{\mathbf{x}}{\text{minimize}} \quad \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \alpha \|\mathbf{x}\|_1. \quad (3)$$

Here, α is the regularization parameter which allows to tune the influence of the two terms in the objective. Naturally, exact recovery of the signal vector cannot be guaranteed in the presence of noise, however, under certain conditions the recovery error can be bounded [10] w.r.t. c_ε .

The discovery of recovery guarantees and error bounds resulted in compressed sensing becoming a very active research field. Many applications have been investigated, among them magnetic resonance imaging (MRI), data acquisition and a multitude of signal processing tasks [11]. A related research field is sparse approximation. While in compressed sensing it is typically assumed that the matrix \mathbf{A} may be modified, in sparse approximation the matrix is fixed and given by the application. Since the recovery algorithms are essentially the same, we do not make a special distinction between these two fields here.

Quite surprising results have been reported when the recovery of non-negative signal vectors, i.e., $\mathbf{x} \in \mathbb{R}_{\geq 0}^N$ is desired. Such non-negativity constraints appear naturally in, e.g., image processing, text & data mining and environmetrics (see also [12]). In the noiseless case, the non-negative signal recovery problem can be formalized as:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \|\mathbf{x}\|_0 \\ & \text{subject to} && \mathbf{b} = \mathbf{A}\mathbf{x} \\ & && \mathbf{x} \geq \mathbf{0}. \end{aligned} \quad (4)$$

Interestingly, the solution space $\{\mathbf{x} \mid \mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$ is a singleton, provided that certain conditions on \mathbf{x} and the matrix \mathbf{A} are fulfilled [13], [14]. Firstly, \mathbf{x} must be sufficiently sparse, where a threshold on the sparsity is defined by the one-sided coherence of the matrix \mathbf{A} . The condition on the matrix states that the row-span of \mathbf{A} must intersect the positive orthant, i.e., a linear combination \mathbf{h} must exist, such that $\mathbf{h}^T \mathbf{A} > \mathbf{0}$.

Let us consider an optimization problem with a general objective function f as well as constraints $\mathbf{A}\mathbf{x} = \mathbf{b}$ and $\mathbf{x} \geq \mathbf{0}$. Furthermore, let the conditions discussed above be fulfilled. Then, since the solution to the linear equation system with non-negativity constraint is unique, for a broad class of objective functions f , including the sparsity inducing ℓ_0 -“norm” and ℓ_1 -norm, it is the solution to the optimization problem as well [13]. This means, that problem (4) can be solved efficiently by a linear program, a result which is similar to the classical compressed sensing scenario.

Turning to the noisy case again, where the measurements are obtained by $\mathbf{b} = \mathbf{A}\mathbf{x} + \varepsilon$, perfect recovery is generally not possible. Similar to the classical compressed sensing case, one could add the non-negativity constraint to problem (3), resulting in the non-negative basis pursuit denoising (NNBPDN) problem. Inspired by the noiseless results above, one could argue that the sparsity inducing objective function may be redundant, so that solving a non-negative least squares (NNLS) problem (see problem (8)) would suffice. This idea was investigated in [15], where subsequent thresholding was suggested to gain truly sparse solutions. In the same work, error bounds were given under the condition that the matrix \mathbf{A} fulfills the self-regularizing property: $\mathbf{h}^T (\frac{1}{M} \mathbf{A}^T \mathbf{A}) \mathbf{h} \geq c_0^2 (\mathbf{1}^T \mathbf{h})^2 \quad \forall \mathbf{h} \geq \mathbf{0}$, where $c_0 \geq 0$ is a universal constant. One advantage of the NNLS approach to solving non-negative sparse approximation problems is that the method is independent of any regularization parameters.

B. (Discontinuous) Dynamical Systems

Dynamical systems with multiple-inputs and multiple-outputs are commonly described in state space. Defining the input vector $\mathbf{u}(t) \in \mathbb{R}^M$, the output vector $\mathbf{y}(t) \in \mathbb{R}^N$ and the state vector $\mathbf{x}(t) \in \mathbb{R}^R$, the dynamical system can be described as

$$\begin{aligned} \dot{\mathbf{x}}(t) &= \mathbf{f}_{\mathbf{x}}(\mathbf{x}(t), \mathbf{u}(t)), \\ \mathbf{y}(t) &= \mathbf{f}_{\mathbf{y}}(\mathbf{x}(t), \mathbf{u}(t)). \end{aligned} \quad (5)$$

Here, $\mathbf{f}_{\mathbf{x}} : \mathbb{R}^R \times \mathbb{R}^M \rightarrow \mathbb{R}^R$ is called the system function and $\mathbf{f}_{\mathbf{y}} : \mathbb{R}^R \times \mathbb{R}^M \rightarrow \mathbb{R}^N$ is called the output function. The term $\dot{\mathbf{x}}(t)$ is the time derivative of the states, i.e., $\dot{\mathbf{x}}(t) = d\mathbf{x}(t)/dt$.

In many cases, it is desirable to analyze dynamical systems for stability. Stable systems ensure a certain level of predictability, thus this concept is very important in control theory. We briefly revise the definitions of stability, which are relevant to this work, following [16]. A point \mathbf{x}_e of a dynamical system with static input $\mathbf{u}(t) = \mathbf{c}_u$ is called *equilibrium point* if

$$\dot{\mathbf{x}}(t) = \mathbf{f}_{\mathbf{x}}(\mathbf{x}_e, \mathbf{u}(t) = \mathbf{c}_u) = \mathbf{0}.$$

We assume $\mathbf{x}_e = \mathbf{0}$ and $\mathbf{c}_u = \mathbf{0}$ without loss of generality in the following since every system can be transformed to this case by a simple linear transformation. An equilibrium point \mathbf{x}_e is called (locally) attractive if all starting points within a neighborhood around the origin $\mathbf{x}_0 \in U(\mathbf{0})$ cause a trajectory $\lim_{t \rightarrow \infty} \mathbf{x}(t) = \mathbf{x}_e = \mathbf{0}$ for the static input $\mathbf{u}(t) = \mathbf{0}$. If this neighborhood is $U(\mathbf{0}) = \mathbb{R}^R$, \mathbf{x}_e is called globally attractive. Additionally, we call an equilibrium point Lyapunov stable if for every ε -neighborhood $U_\varepsilon(\mathbf{0}) = \{\gamma \in \mathbb{R}^R \mid \|\gamma\|_2 < \varepsilon\}$, there is a δ -neighborhood $U_\delta(\mathbf{0}) = \{\gamma \in \mathbb{R}^R \mid \|\gamma\|_2 < \delta\}$ such that if $\mathbf{x}_0 \in U_\delta(\mathbf{0})$, then $\mathbf{x}(t) \in U_\varepsilon(\mathbf{0})$ for every $t > 0$ and $\delta = \delta(\varepsilon)$. Finally, if the equilibrium is locally (globally) attractive and Lyapunov stable, it is called locally (globally) asymptotically stable. In summary, if asymptotic stability holds, the dynamical system with constant input and a starting point \mathbf{x}_0 close enough to the equilibrium point will generate a trajectory converging to \mathbf{x}_e while the trajectory is guaranteed to stay within a certain range of \mathbf{x}_e in the process.

For *classical* dynamical systems, when analyzing stability there is the restriction that the system function $\mathbf{f}_{\mathbf{x}}$ must be continuously differentiable [17] or more generally Lipschitz continuous [18], [19]. This guarantees that a solution $\mathbf{x}(t)$ to the differential vector equation (5) exists and that it is unique for a starting point \mathbf{x}_0 . When studying *discontinuous* dynamical systems, this restriction is dropped so that $\mathbf{f}_{\mathbf{x}}$ can be discontinuous. Stability analysis of such systems is more complicated, since existence and uniqueness of solutions is not guaranteed anymore. Also, there are different kinds of generalized solutions which have been studied, such as Caratheodory or Filippov solutions [20]. In this work, we make use of a Caratheodory solution. In simple terms, such a solution generally satisfies the system equations, except for a set of time instances of measure zero (e.g., a countable, not necessarily finite, set). Formally, a Caratheodory solution can be defined using the Lebesgue integral as an absolutely continuous function $\mathbf{x}(t)$, which satisfies [20]

$$\mathbf{x}(t) = \mathbf{x}_0 + \int_0^t \mathbf{f}_{\mathbf{x}}(\mathbf{x}(t'), \mathbf{u}(t')) dt'.$$

III. SYSTEM MODEL

We consider a continuously valued neural network with input vector $\mathbf{u}(t) \in \mathbb{R}^M$ and non-negative output vector $\mathbf{y}(t) \in \mathbb{R}_{\geq 0}^N$. For the upcoming analysis, we represent the network as a discontinuous system in state space with state vector $\mathbf{x}(t)$ (see also Figure 1). From this, it follows that the matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$. The integrator denoted by \int^+ is a non-linear integrator as defined in the following. If all states are positive, this block behaves like a standard integrator, however,




Fig. 1. Block diagram of the discontinuous system with a non-linear integrator.

if a state reaches zero, integration of this state is halted. Thus, said state remains at value zero until it is subjected to a positive gradient again. Defining the non-linear integrator as $\mathbf{x}(t) = \int_0^t \dot{\mathbf{x}}(\tau) d\tau$, it follows that the derivative of the state vector $\dot{\mathbf{x}}(t)$ must be discontinuous:

$$\dot{x}_i(t) = \begin{cases} \tilde{x}_i(t), & \text{if } x_i(t) > 0 \\ [\tilde{x}_i(t)]^+, & \text{if } x_i(t) = 0 \\ c, & \text{if } x_i(t) < 0. \end{cases} \quad i = 1, \dots, N. \quad (6)$$

Here, $\tilde{x}_i(t)$ is the input of the i -th integrator and the positive part of the i -th vector element is denoted as $[\tilde{x}_i]^+ = \max\{0, \tilde{x}_i\}$. The inclusion of the constant $c > 0$ is a technicality to ensure recovery from non-negative states due to faulty initial conditions or disturbances. See Figure 2 for an example of the behavior of the integrator described in (6).




Fig. 2. Comparison between the output signals of a standard linear integrator (dashed blue line) and a non-linear integrator as described above (solid blue line). The input signal (dotted green line) is plotted according to the second ordinate on the right.

Using (6) and

$$\tilde{\mathbf{x}}(t) = \mathbf{A}^T \mathbf{u}(t) - \mathbf{A}^T \mathbf{A} \mathbf{x}(t),$$

the system equations can be formalized for the discontinuous system depicted in Figure 1.

$$\begin{aligned} \dot{\mathbf{x}}(t) &= \mathbf{f}_{\mathbf{x}}(\mathbf{x}(t), \mathbf{u}(t)), \quad \text{with } \dot{x}_i(t) = f_{x_i}(\mathbf{x}(t), \mathbf{u}(t)), \\ \dot{x}_i(t) &= \begin{cases} [\tilde{\mathbf{x}}(t)]_i, & \text{if } x_i(t) > 0 \\ [\tilde{\mathbf{x}}(t)]_i^+, & \text{if } x_i(t) = 0 \\ c, & \text{if } x_i(t) < 0 \end{cases} \quad i = 1, \dots, N \\ \mathbf{y}(t) &= \mathbf{f}_{\mathbf{y}}(\mathbf{x}(t), \mathbf{u}(t)) = \mathbf{x}(t). \end{aligned} \quad (7)$$

The system defined by (7) has a very interesting property: it is able to solve non-negative least squares (NNLS) optimization problems. To further clarify this feature, consider the following optimization problem with optimal point \mathbf{x}^* and $\mathbf{b} \in \mathbb{R}^M$:

$$\begin{aligned} &\underset{\mathbf{x}}{\text{minimize}} && \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 \\ &\text{subject to} && \mathbf{x} \geq \mathbf{0}. \end{aligned} \quad (8)$$

This is a convex optimization problem, which does not possess an analytical solution, in contrast to an unconstrained least-squares problem [21].

If the discontinuous system of Figure 1 receives the constant input $\mathbf{u}(t) = \mathbf{b}$, it will converge to its equilibrium point $\mathbf{x}_e = \mathbf{x}^*$, which is also the optimal point of the NNLS optimization problem (8). We formally prove this fact in Section IV-A.

IV. CONVERGENCE ANALYSIS

In this section, we take a closer look at the convergence behavior of the discontinuous dynamic system shown in Figure 1. First, we show that the equilibrium points of the system are solutions of NNLS optimization problems in Section IV-A. Second, we characterize the Caratheodory solution of the differential equation in Section IV-B, which is a prerequisite for the stability analysis. Then, we turn to the actual stability analysis in Section IV-C, which is based on Lyapunov's direct method. Later, we additionally prove in Section IV-D that the number of switches performed by the non-linear integrator during convergence to the equilibrium point is finite under a mild condition on the solution. Finally, we briefly discuss in Section IV-E how the convergence analysis can be generalized to limited integrators with box constraints of the form $\mathbf{l} \leq \mathbf{x}(t) \leq \mathbf{h}$.

Before we begin the convergence analysis, we introduce some notation which will be helpful in subsequent sections of this text. First, we denote the set of all states' indices as $\mathcal{I} = \{1, \dots, N\}$. Then, we divide this set into three time-varying sets $\mathcal{I}^+(t)$, $\mathcal{I}^0(t)$ and $\mathcal{I}^-(t)$. These can informally be described as the indices of the system states that are positive, zero and negative at time t , respectively. A formal definition of the sets can be given as follows:

$$\begin{aligned} \mathcal{I}^+(t) &= \left\{ i \in \mathcal{I} \mid x_i(t) > 0 \right\} \\ &\cup \left\{ i \in \mathcal{I} \mid x_i(t) = 0, \tilde{x}_i(t) \geq 0 \right\}, \\ \mathcal{I}^0(t) &= \left\{ i \in \mathcal{I} \mid x_i(t) = 0, \tilde{x}_i(t) < 0 \right\}, \\ \mathcal{I}^-(t) &= \left\{ i \in \mathcal{I} \mid x_i(t) < 0 \right\}. \end{aligned} \quad (9)$$

Note, that the set $\mathcal{I}^+(t)$ also contains indices of states which are zero at time t if its corresponding integrator input is positive. Also, no state index can be present in two or all of the sets $\mathcal{I}^+(t)$, $\mathcal{I}^0(t)$ and $\mathcal{I}^-(t)$ at any given time t . Keep in mind, that states may only appear in $\mathcal{I}^-(t)$ if faulty initial conditions are present or disturbances cause a state to become negative. Otherwise, it is impossible for a state to become negative.

We use the presented notation to create vectors from a selection of entries from another vector and to create matrices, which are a selection of a set of rows or columns from another matrix. For example, a vector containing all states whose indices' are given in the set \mathcal{I}^+ is denoted as $[\mathbf{x}(t)]_{\mathcal{I}^+}$. Similarly, a matrix containing only the columns specified by the indices in the set \mathcal{I}^+ is written as $[\mathbf{A}]_{:, \mathcal{I}^+}$.

A. Equilibrium Point

First, we prove that the equilibrium point of the dynamic system depicted in Figure 1 is a solution of the NNLS optimization problem (8). To that end, we first review the KKT conditions of the NNLS problem. Let $f_0(\mathbf{x})$ refer to its objective function and $f_i(\mathbf{x}) = -x_i \leq 0$ for $i = 1, \dots, N$ to its inequality constraints in standard form. Their gradients are $\nabla f_0(\mathbf{x}) = \mathbf{A}^T \mathbf{A} \mathbf{x} - \mathbf{A}^T \mathbf{b}$ and $\nabla f_i(\mathbf{x}) = -\mathbf{e}_i$ for $i = 1, \dots, N$, respectively. Here, \mathbf{e}_i is the i -th standard basis vector. From this, we can derive the KKT conditions of problem (8) as:

$$\begin{aligned} \mathbf{x}^* &\geq \mathbf{0} \\ \boldsymbol{\lambda}^* &\geq \mathbf{0} \\ \lambda_i^* x_i^* &= 0, \quad i = 1, \dots, N \\ \mathbf{A}^T \mathbf{A} \mathbf{x}^* - \mathbf{A}^T \mathbf{b} - \boldsymbol{\lambda}^* &= \mathbf{0}. \end{aligned} \quad (10)$$

Obviously, there exists a strictly feasible point, i.e., $\mathbf{x} > \mathbf{0}$. Thus, Slater's condition holds for problem (8). As a consequence, strong duality holds and the KKT conditions presented in (10) are necessary and sufficient conditions for optimality [21]. Therefore, the points \mathbf{x}^* and $\boldsymbol{\lambda}^*$ are primal and dual optimal. If the matrix \mathbf{A} has full column-rank, the solution of problem (8) is unique [22].

For an equilibrium point \mathbf{x}_e of a dynamic system with static input $\mathbf{u}(t) = \mathbf{u}_e$

$$\dot{\mathbf{x}}(t) = \mathbf{f}_{\mathbf{x}}(\mathbf{x}_e, \mathbf{u}_e) = \mathbf{0}. \quad (11)$$

must hold [16]. In other words, the system states stay the same for constant input \mathbf{u}_e , since the time derivatives of the system states vanish at the equilibrium point \mathbf{x}_e .

Let the input of the system described by (7) be constant $\mathbf{u}(t) = \mathbf{b}$ and define

$$\boldsymbol{\lambda}_e := \mathbf{A}^T \mathbf{A} \mathbf{x}_e - \mathbf{A}^T \mathbf{b}. \quad (12)$$

As stated by (11), for an equilibrium point \mathbf{x}_e it must hold for $i = 1, \dots, N$:

$$0 = \dot{x}_i(t) = f_{x_i}(\mathbf{x}_e, \mathbf{b}) = \begin{cases} [-\boldsymbol{\lambda}_e]_i, & \text{if } [\mathbf{x}_e]_i > 0 \\ [-\boldsymbol{\lambda}_e]_i^+, & \text{if } [\mathbf{x}_e]_i = 0 \\ c, & \text{if } [\mathbf{x}_e]_i < 0 \end{cases} \quad (13)$$

It can be easily seen that no $[\mathbf{x}_e]_i$ can be negative since $c > 0$. Thus, $\mathbf{x}_e \geq \mathbf{0}$ must hold. For the first case of (13), i.e., $[\mathbf{x}_e]_i > 0$, it follows that $[\boldsymbol{\lambda}_e]_i = 0$. Similarly, for the second case of (13), i.e., $[\mathbf{x}_e]_i = 0$, we must have $[\boldsymbol{\lambda}_e]_i \geq 0$. In summary, we can write these conditions as:

$$\begin{aligned} \mathbf{x}_e &\geq \mathbf{0} \\ \boldsymbol{\lambda}_e &\geq \mathbf{0} \\ [\boldsymbol{\lambda}_e]_i [\mathbf{x}_e]_i &= 0, \quad i = 1, \dots, N \\ \mathbf{A}^T \mathbf{A} \mathbf{x}_e - \mathbf{A}^T \mathbf{b} - \boldsymbol{\lambda}_e &= \mathbf{0}. \end{aligned} \quad (14)$$

Comparing (10) and (14), we see that $(\mathbf{x}_e, \boldsymbol{\lambda}_e)$ satisfy the KKT conditions of the NNLS optimization problem and therefore are a primal and dual optimal solution. This means that given the constant system input $\mathbf{u}(t) = \mathbf{b}$, the discontinuous system depicted in Figure 1 has equilibrium points that are in fact the solution of the NNLS problem as defined in (8).

B. Caratheodory Solution

First, we reformulate system (7) into the following equivalent formulation:

$$\dot{x}_i(t) = \begin{cases} \tilde{x}_i(t), & \begin{cases} \text{if } x_i(t) > 0 \\ \text{if } x_i(t) = 0, \tilde{x}_i(t) \geq 0 \end{cases} \\ 0, & \text{if } x_i(t) = 0, \tilde{x}_i(t) < 0 \\ c, & \text{if } x_i(t) < 0. \end{cases} \quad (15)$$

Here, the positive part present in the second case in (7) was manually split in (15). Note also that the cases of (15) correspond to the definitions of the index sets in (9).

Let us investigate a time span, in which the index sets from (9) are constant, i.e., time instances τ in an interval $\tau \in [t_j, t_{j+1})$ with $\mathcal{I}^+(\tau) = \mathcal{I}_j^+$, $\mathcal{I}^0(\tau) = \mathcal{I}_j^0$ and $\mathcal{I}^-(\tau) = \mathcal{I}_j^-$. For such a time span, the system (15) can be split into several linear subsystems. States in \mathcal{I}_j^0 or \mathcal{I}_j^- are not influenced by other states, hence forming independent linear subsystems. A new linear subsystem can be formed for the states in \mathcal{I}_j^+ , which is independent of the states in \mathcal{I}_j^0 and the influence from the states in \mathcal{I}_j^- can be counted as additional input, since the states in \mathcal{I}_j^- are independent from others. Defining

$$\begin{aligned} \mathbf{x}_j^+(\tau) &= [\mathbf{x}(\tau)]_{\mathcal{I}_j^+}, \\ \mathbf{A}_j^+ &= [\mathbf{A}]_{\mathcal{I}_j^+}, \\ \mathbf{u}_j^+(\tau) &= \mathbf{u}(\tau) - [\mathbf{A}]_{\mathcal{I}_j^-} [\mathbf{x}(\tau)]_{\mathcal{I}_j^-}, \end{aligned} \quad (16)$$

the system differential equation for this subsystem can be given as

$$\dot{\mathbf{x}}_j^+(\tau) = \mathbf{A}_j^+ \mathbf{u}_j^+(\tau) - \mathbf{A}_j^{+T} \mathbf{A}_j^+ \mathbf{x}_j^+(\tau). \quad (17)$$

Thus, provided the matrix \mathbf{A}_j^+ has full column-rank, this subsystem is globally asymptotically stable. This follows since the subsystem is linear and its system matrix $-\mathbf{A}_j^{+T} \mathbf{A}_j^+$ is negative definite, i.e., all of its eigenvalues are negative [23].

Next, we inspect the transitions of states between the index sets (9). The index j , already used in the preceding paragraph, identifies these transitions and can be used to specify time intervals $[t_j, t_{j+1})$ with constant index sets. Assume $\tau \in [t_j, t_{j+1})$ in the following:

- $\mathcal{I}_j^+ \xrightarrow{i} \mathcal{I}_{j+1}^0$: When a state $i \in \mathcal{I}_j^+$ in the subsystem (17) reaches zero, it transits from \mathcal{I}_j^+ to \mathcal{I}_{j+1}^0 . When reaching zero, the states' influence vanishes. Consequently, this transition does not cause discontinuities in any other states' trajectories.
- $\mathcal{I}_j^0 \xrightarrow{i} \mathcal{I}_{j+1}^+$: If the input $\mathbf{u}(\tau)$ or changes in the values of states in \mathcal{I}_j^+ or \mathcal{I}_j^- cause the term $\tilde{x}_i(\tau)$ to become greater or equal to zero, the state i transits from \mathcal{I}_j^0 to \mathcal{I}_{j+1}^+ . However, since its immediate influence on the other states is zero, this transition is continuous.
- $\mathcal{I}_j^- \xrightarrow{i} \mathcal{I}_{j+1}^0$ or $\mathcal{I}_j^- \xrightarrow{i} \mathcal{I}_{j+1}^+$: At time t_j , the state $x_i(t_j)$ must be negative, thus at constant rate of integration it will become zero at $t_{j+1} = t_j - \frac{x_i(t_j)}{c}$. Depending on $\tilde{x}_i(t_{j+1})$ it will transit to either \mathcal{I}_{j+1}^0 or \mathcal{I}_{j+1}^+ . Note, that this transition is also continuous, since the immediate influence of the state i is zero.
- $\mathcal{I}_j^+ \xrightarrow{i} \mathcal{I}_{j+1}^-$ or $\mathcal{I}_j^0 \xrightarrow{i} \mathcal{I}_{j+1}^-$: These transitions are impossible with the system definition (15).

In summary, all transitions are continuous. Provided that the system matrix \mathbf{A} has full column-rank, the behavior of the subsystem (17) is unique at every point in time. Also, the behavior of the subsystems formed by \mathcal{I}_j^0 and \mathcal{I}_j^- is unique. Finally, the transitions are obviously unique as well and thus it follows the solution $\mathbf{x}(t)$ must be unique. Let $\mathcal{T}(t)$ be the set containing all transition times up until time t , i.e., $\mathcal{T}(t) = \{t_j \mid t_j \leq t\}$. Note, that $\mathcal{T}(t)$ is a countable set with measure zero. Defining $t_0 := 0$ and $t_{|\mathcal{T}(t)|+1} := t$, the unique Caratheodory solution of the system follows as:

$$\mathbf{x}(t) = \mathbf{x}_0 + \sum_{j=0}^{|\mathcal{T}(t)|} \int_{t_j}^{t_{j+1}} \mathbf{f}_{\mathbf{x}}(\mathbf{x}(t'), \mathbf{u}(t')) dt'. \quad (18)$$

The individual integrals in (18) can be given using the notation of the subsystems in (16) for $j = 1, \dots, |\mathcal{T}(t)|$ as:

$$\begin{aligned} \left[\int_{t_j}^{t_{j+1}} \dots dt' \right]_{\mathcal{I}_j^+} &= \int_{t_j}^{t_{j+1}} e^{-\mathbf{A}_j^+ T \mathbf{A}_j^+ (t_{j+1}-t')} \mathbf{A}_j^+ \mathbf{u}_j^+(t') dt', \\ \left[\int_{t_j}^{t_{j+1}} \dots dt' \right]_{\mathcal{I}_j^0} &= \mathbf{0}, \\ \left[\int_{t_j}^{t_{j+1}} \dots dt' \right]_{\mathcal{I}_j^-} &= c(t_{j+1} - t_j) \mathbf{1}. \end{aligned} \quad (19)$$

Thus, the Caratheodory solution is sufficiently defined. Finding a closed form of (18), however, requires expressions for the transition times and a closed form of the first equation in (19), which can be very challenging.

C. Stability

In order to show asymptotic stability of the equilibrium point \mathbf{x}_e for the system in (7), we apply a linear transformation as discussed in Section II-B. We substitute $\mathbf{z}(t) = \mathbf{x}(t) - \mathbf{x}_e$ and $\hat{\mathbf{u}}(t) = \mathbf{u}(t) - \mathbf{b}$, so that the equilibrium point of the transformed system is $\mathbf{z}_e = \mathbf{0}$ for the constant input $\hat{\mathbf{u}}(t) = \mathbf{0}$. The transformed system function can be found using:

$$\dot{\mathbf{z}}(t) = \mathbf{f}_{\mathbf{z}}(\mathbf{z}(t), \hat{\mathbf{u}}(t)) = \mathbf{f}_{\mathbf{x}}(\mathbf{z}(t) + \mathbf{x}_e, \hat{\mathbf{u}}(t) + \mathbf{b}).$$

Note that we omit the time index t in the following derivations to increase readability. Using (12), splitting the cases and simplifying, we arrive at:

$$\begin{aligned} \dot{\mathbf{z}}_i &= \begin{cases} [-\mathbf{A}^T \mathbf{A} \mathbf{z} - \mathbf{A}^T \mathbf{A} \mathbf{x}_e + \mathbf{A}^T \mathbf{b}]_i, & \text{if } z_i + [\mathbf{x}_e]_i > 0 \\ [-\mathbf{A}^T \mathbf{A} \mathbf{z} - \mathbf{A}^T \mathbf{A} \mathbf{x}_e + \mathbf{A}^T \mathbf{b}]_i^+, & \text{if } z_i + [\mathbf{x}_e]_i = 0 \\ c, & \text{if } z_i + [\mathbf{x}_e]_i < 0 \end{cases} \\ &= \begin{cases} [-\mathbf{A}^T \mathbf{A} \mathbf{z}]_i, & \text{if } z_i + [\mathbf{x}_e]_i > 0, [\mathbf{x}_e]_i > 0 \\ [-\mathbf{A}^T \mathbf{A} \mathbf{z} - \lambda_e]_i, & \text{if } z_i > 0, [\mathbf{x}_e]_i = 0 \\ [-\mathbf{A}^T \mathbf{A} \mathbf{z}]_i^+, & \text{if } z_i + [\mathbf{x}_e]_i = 0, [\mathbf{x}_e]_i > 0 \\ [-\mathbf{A}^T \mathbf{A} \mathbf{z} - \lambda_e]_i^+, & \text{if } z_i = 0, [\mathbf{x}_e]_i = 0 \\ c, & \text{if } z_i + [\mathbf{x}_e]_i < 0 \end{cases} \end{aligned} \quad (20)$$

In the following, we study stability using Lyapunov's direct method. For discontinuous dynamical systems, there exists a lot of dedicated work on stability theory, e.g., for systems

with Caratheodory solutions. However, if a smooth Lyapunov function is applicable, the theory reduces to classical stability analysis [20]. Indeed, a smooth Lyapunov function suffices for this system and it is chosen as:

$$V(\mathbf{z}) = \frac{1}{2} \mathbf{z}^T \mathbf{z},$$

with gradient $\nabla V(\mathbf{z}) = \mathbf{z}$. The time derivative of $V(\mathbf{z})$ yields

$$\dot{V}(\mathbf{z}) = [\nabla V(\mathbf{z})]^T \dot{\mathbf{z}} = \mathbf{z}^T \dot{\mathbf{z}} = \sum_{i=1}^N z_i \dot{z}_i.$$

Three conditions must be fulfilled such that the chosen function is a Lyapunov function for the system and asymptotic stability holds:

$$\begin{aligned} V(\mathbf{0}) &= 0, \\ V(\mathbf{z}) &= \frac{1}{2} \|\mathbf{z}\|_2^2 > 0 \quad \forall \mathbf{z} \in \mathbb{R}^N \setminus \{\mathbf{0}\}, \\ \dot{V}(\mathbf{z}) &< 0 \quad \forall \mathbf{z} \in \mathbb{R}^N \setminus \{\mathbf{0}\}. \end{aligned} \quad (21)$$

Obviously, the first two conditions are fulfilled. With (20), we can write

$$\begin{aligned} z_i \dot{z}_i &= \begin{cases} z_i [-\mathbf{A}^T \mathbf{A} \mathbf{z}]_i, & \text{if } \begin{cases} z_i + [\mathbf{x}_e]_i > 0, \\ [\mathbf{x}_e]_i > 0 \end{cases} \\ z_i [-\mathbf{A}^T \mathbf{A} \mathbf{z} - \lambda_e]_i, & \text{if } \begin{cases} z_i > 0, \\ [\mathbf{x}_e]_i = 0 \end{cases} \\ z_i [-\mathbf{A}^T \mathbf{A} \mathbf{z}]_i^+, & \text{if } \begin{cases} z_i + [\mathbf{x}_e]_i = 0, \\ [\mathbf{x}_e]_i > 0 \end{cases} \\ z_i [-\mathbf{A}^T \mathbf{A} \mathbf{z} - \lambda_e]_i^+, & \text{if } \begin{cases} z_i = 0, \\ [\mathbf{x}_e]_i = 0 \end{cases} \\ z_i c, & \text{if } z_i + [\mathbf{x}_e]_i < 0. \end{cases} \end{aligned} \quad \begin{aligned} (22a) \\ (22b) \\ (22c) \\ (22d) \\ (22e) \end{aligned}$$

Note that it was shown in Section IV-A that $\lambda_e \geq 0$. For $\mathbf{z} \in U_\varepsilon(\mathbf{0})$, it follows for the subcases:

$$\begin{aligned} (22b) &= -z_i [\mathbf{A}^T \mathbf{A} \mathbf{z}]_i - \underbrace{z_i}_{\geq 0} \underbrace{[\lambda_e]_i}_{\geq 0} \leq -z_i [\mathbf{A}^T \mathbf{A} \mathbf{z}]_i \\ (22c) &= \underbrace{z_i}_{= -[\mathbf{x}_e]_i \leq 0} \underbrace{[\mathbf{A}^T \mathbf{A} \mathbf{z}]_i^+}_{\geq [\mathbf{A}^T \mathbf{A} \mathbf{z}]_i} \leq -z_i [\mathbf{A}^T \mathbf{A} \mathbf{z}]_i \\ (22d) &= \underbrace{z_i}_{= 0} [-\mathbf{A}^T \mathbf{A} \mathbf{z} - \lambda_e]_i^+ = 0 = -\underbrace{z_i}_{= 0} [\mathbf{A}^T \mathbf{A} \mathbf{z}]_i \\ (22e) &= \underbrace{z_i}_{\leq 0} \underbrace{c}_{\geq -[\mathbf{A}^T \mathbf{A} \mathbf{z}]_i} \leq -z_i [\mathbf{A}^T \mathbf{A} \mathbf{z}]_i. \end{aligned}$$

Thus, for $\mathbf{z} \in \mathbb{R}^N \setminus \{\mathbf{0}\}$ and \mathbf{A} having full column-rank, it holds (almost everywhere):

$$\begin{aligned} \dot{V}(\mathbf{z}) &= \sum_{i=1}^N z_i \dot{z}_i \leq \sum_{i=1}^N -z_i [\mathbf{A}^T \mathbf{A} \mathbf{z}]_i \\ &= -\mathbf{z}^T \mathbf{A}^T \mathbf{A} \mathbf{z} < 0. \end{aligned} \quad (23)$$

In the last inequality, we used the fact that $\mathbf{A}^T \mathbf{A}$ is a positive definite matrix if \mathbf{A} has full column-rank. The function $V(\mathbf{z})$ satisfies the conditions (21) and hence the equilibrium point $\mathbf{z} = \mathbf{0}$ is asymptotically stable in $U_\varepsilon(\mathbf{0})$. The parameter ε present in the neighborhood $U_\varepsilon(\mathbf{0})$ can be chosen arbitrarily, therefore semi-global stability holds for the discontinuous system (7).

D. Finite Number of Switches

In Section IV-B, we already established that the set of transition times $\mathcal{T}(t)$ is countable. Hence, if there is a time after which no more switches (transitions) occur, then the number of switches is finite. We show that under a mild condition on the solution, this is indeed the case. For the following analysis, we define the time-dependent vector

$$\lambda(t) = -\tilde{x}(t) = \mathbf{A}^T \mathbf{A}x(t) - \mathbf{A}^T \mathbf{b},$$

which is the negative integrator input. Above, it was already shown that the equilibrium point \mathbf{x}_e is unique if \mathbf{A} has full column-rank. Similar to (9), we can define unique index sets for the equilibrium point as:

$$\begin{aligned}\mathcal{I}_e^+ &= \left\{ i \in \mathcal{I} \mid [\mathbf{x}_e]_i > 0 \right\} \\ &\cup \left\{ i \in \mathcal{I} \mid [\mathbf{x}_e]_i = 0, -[\lambda_e]_i \geq 0 \right\}, \\ \mathcal{I}_e^0 &= \left\{ i \in \mathcal{I} \mid [\mathbf{x}_e]_i = 0, -[\lambda_e]_i < 0 \right\}, \\ \mathcal{I}_e^- &= \emptyset.\end{aligned}$$

To prevent infinite oscillations at the transition boundaries, we assume there exists $\varepsilon_x, \varepsilon_\lambda > 0$, such that

$$\begin{aligned}\forall i \in \mathcal{I}_e^+ : \quad & [\mathbf{x}_e]_i > \varepsilon_x \\ \forall i \in \mathcal{I}_e^0 : \quad & [\lambda_e]_i > 2\varepsilon_\lambda.\end{aligned}$$

Essentially, this means the equilibrium point \mathbf{x}_e is good mannered in the sense that it is not directly on a transition boundary (see also [9]). Note that under this assumption also

$$\begin{aligned}\forall i \in \mathcal{I}_e^+ : \quad & [\lambda_e]_i = 0 \\ \forall i \in \mathcal{I}_e^0 : \quad & [\mathbf{x}_e]_i = 0\end{aligned}$$

holds.

The system states $\mathbf{x}(t)$ converge to the equilibrium point \mathbf{x}_e and likewise $\lambda(t)$ converges to λ_e . Additionally, due to the behavior of the bounded integrator (6), after a limited time all states must be non-negative. Thus, there exists a time $T > 0$ so that $\forall t \geq T$:

$$\begin{aligned}\|\mathbf{x}(t) - \mathbf{x}_e\|_2 &< \varepsilon_x \\ \|\lambda(t) - \lambda_e\|_2 &< \varepsilon_\lambda \\ \mathbf{x}(t) &\geq \mathbf{0}.\end{aligned}$$

Similar to the approach in [9], we can say that for $\forall t \geq T$ and $\forall i \in \mathcal{I}_e^+$:

$$\begin{aligned}\varepsilon_x &> |x_i - [\mathbf{x}_e]_i| \geq |[\mathbf{x}_e]_i| - |x_i| > \varepsilon_x - |x_i(t)| \\ &\Rightarrow |x_i(t)| > 0 \\ &\Rightarrow x_i(t) > 0.\end{aligned}$$

As can be seen, for all times $t \geq T$ the states in \mathcal{I}_e^+ stay strictly positive and hence do not switch anymore. In other words, all states in \mathcal{I}_e^+ have a certain gap to zero, that is $[\mathbf{x}_e]_{\mathcal{I}_e^+} > \varepsilon_x$ and after the time T all states in \mathcal{I}_e^+ stay closer to the equilibrium value than this gap.

Likewise, for $t \geq T$ and $\forall i \in \mathcal{I}_e^0$ it follows:

$$\begin{aligned}\varepsilon_\lambda &> |\lambda_i - [\lambda_e]_i| \geq |[\lambda_e]_i| - |\lambda_i| > 2\varepsilon_\lambda - |\lambda_i(t)| \\ &\Rightarrow |\lambda_i(t)| > \varepsilon_\lambda > 0 \\ &\Rightarrow \lambda_i(t) > \varepsilon_\lambda > 0.\end{aligned}$$

So, for all times $t \geq T$, the integrator inputs $\tilde{x}_i(t) = -\lambda_i(t)$ of the states $i \in \mathcal{I}_e^0$ stay strictly negative. At time T , however, some states in \mathcal{I}^0 may still be positive and some may be zero already. We investigate both cases separately:

- zero states, i.e., $\forall i \in \mathcal{I}_e^0 \cap \mathcal{I}^0(T)$: As explained above, the integrator input stays negative. Thus, no transitions occur after time T .
- non-zero states, i.e., $\forall i \in \mathcal{I}_e^0 \cap \mathcal{I}^+(T)$: For $t \geq T$, it holds

$$\begin{aligned}x_i(t) &= x_i(T) + \int_T^t -\lambda_i(t') dt' \\ &< x_i(T) + \int_T^t -\varepsilon_\lambda dt' = x_i(T) - \varepsilon_\lambda(t - T).\end{aligned}\tag{24}$$

Thus, a state i reaches zero at time $\frac{x_i(T)}{\varepsilon_\lambda} + T$ at the latest and stays zero since the integrator input remains negative.

Hence, after time

$$\max_{i \in \mathcal{I}_e^0 \cap \mathcal{I}^+(T)} \frac{x_i(T)}{\varepsilon_\lambda} + T,$$

no more switches occur. Since this point in time is finite, the number of switches is finite as well.

E. Generalization to Box Constraints

In this section, we briefly discuss an extension to system (7), so that it is able to solve optimization problems of the form

$$\begin{aligned}\underset{\mathbf{x}}{\text{minimize}} \quad & \frac{1}{2} \mathbf{x}^T \mathbf{P} \mathbf{x} - \mathbf{q}^T \mathbf{x} \\ \text{subject to} \quad & \mathbf{l} \leq \mathbf{x} \leq \mathbf{h}.\end{aligned}\tag{25}$$

If the matrix \mathbf{P} is positive definite, which we assume in the following, this is a convex quadratic program (QP) with lower and upper bounds, see e.g. [21]. By defining $\mathbf{P} = \mathbf{A}^T \mathbf{A}$ and $\mathbf{q} = \mathbf{A}^T \mathbf{b}$ this problem is identical to a bounded least squares problem.

The integrator used in system (7) is replaced by an integrator with general lower and upper bounds defined by \mathbf{l} and \mathbf{h} , respectively. Also, the system matrix is replaced by the matrix \mathbf{P} . Defining the new integrator input as $\hat{\mathbf{x}}(t) = \mathbf{u}(t) - \mathbf{P}\mathbf{x}(t)$, we can give the system equations of the modified system as:

$$\dot{x}_i(t) = \begin{cases} -c, & \text{if } x_i(t) > h_i \\ [\hat{\mathbf{x}}(t)]_i^-, & \text{if } x_i(t) = h_i \\ [\hat{\mathbf{x}}(t)]_i, & \text{if } l_i < x_i(t) < h_i \\ [\hat{\mathbf{x}}(t)]_i^+, & \text{if } x_i(t) = l_i \\ c, & \text{if } x_i(t) < l_i \end{cases} \quad i = 1, \dots, N\tag{26}$$

$$\mathbf{y}(t) = \mathbf{x}(t).$$

If system (26) receives the constant input $\mathbf{u}(t) = \mathbf{q}$, it will converge to its equilibrium point, which is the solution of the QP (25). This can be shown by using the definition of an equilibrium point and comparing it to the KKT conditions of problem (25). To show that a Caratheodory solution exists, one can reformulate the system equations similar to (15). Corresponding to the cases present in the system function, five index sets (compare (9)) can be defined: one for states

which are strictly inside the bounds forming an independent linear system, two for states at the bounds and finally two for states which exceed the bounds due to initial conditions or disturbances. Transitions can be shown to be smooth and the Caratheodory solution can be characterized. Asymptotic stability holds as well. After linearly transforming the system, the same Lyapunov function as above may be used. Similarly to (22), 11 cases may be worked out and the Lyapunov function can be shown to be upper bounded by $-z^T P z < 0$, as we assumed P to be positive definite.

V. APPLICATION: OLFACTORY MIXTURE DECOMPOSITION

The first application of the dynamic system (7) we present is the decomposition of a mixture of sensory stimuli as a model for a biological neural network. It is based on the structure of the olfactory system in mammals or more precisely the main olfactory bulb (MOB). We first present the model and discuss numerical results thereafter.

A. Model Description

The olfactory system is considered a flat system since it consists of few processing stages before olfactory information is transmitted to higher emotional or cognitive brain regions. Olfactory sensory neurons (OSNs) are the chemical sensors of the system. By expressing different olfactory receptor genes, there are approximately 1000 “types” of OSNs. There are around 1000 to 10000 OSNs of each type in mice / rodents, resulting in a total number of OSNs between 10^6 and 10^7 , see [24]–[27]. In the main olfactory bulb, all OSNs of the same type converge into two spherical neuropil structures of high connectivity, called glomeruli. In turn, connected to one glomerulus are approximately 25 mitral cells (MCs) [24], which finally project to higher brain areas. These parallel processing paths are influenced by the periglomerular cells (PGCs) and interact with each other via short axon cells (SACs) on the glomeruli level and through granule cells (GCs) on the mitral cell level, compare also Figure 3. It is believed, these so called lateral interactions play a key role in the signal processing within the MOB, see e.g. [25], [27], [28]. Despite the MOB’s biological significance, however, the principles behind its information processing still remain unclear.

The $MC \rightarrow GC \rightarrow MC$ interconnection circuit seems to be unique in the body and has a peculiar behavior. An excited MC that is connected to a GC will be inhibited via the connection (self-inhibition). Another MC, which is also connected to the same GC, will be inhibited as well. So in total the GC processing is inhibitory and thus lowers the activity of connected MCs.

Nerve cells are modeled by a dimensionless positive activity level which expresses their level of excitation. This way, we can abstract from the electro-chemical level of action potentials and do not need to make additional assumptions about the existence of rate- or timing codes, while still retaining a proportional relationship. Furthermore, the GCs are modeled as passive components and the MCs are the only active components in the system. While the GCs are able to

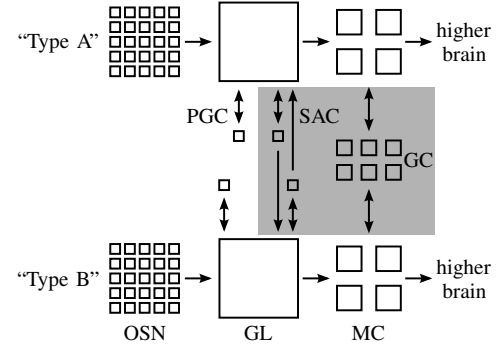


Fig. 3. Highly abstracted block diagram of the main olfactory bulb. Two parallel pathways, associated with two types (olfactory receptor genes) of OSNs, are shown. The sources of potential lateral interactions (SACs and GCs) are depicted within the gray shaded area. Also, the number of cells involved for each cell type is qualitatively reflected.

generate action potentials, they seem to do so only rarely and may thus be covered in a passive manner by the model.

Let all the GCs and MCs in the system be identifiable by indices. Then we can model the activity level of a GC with index j at a given time t as follows:

$$l_j(t) = \sum_i w_{ji} x_i(t), \quad (27)$$

where x_i is the output of the MC with index i and $w_{ji} \geq 0$ is a weight describing the synaptic connection. We assume for this model that multiple GCs connected to one MC are only present to ease synaptic routing in the body. Thus in our model, through one GC a MC can be connected to all other MCs with the corresponding weights. Likewise, the i -th MC can be modeled as:

$$\begin{aligned} \tilde{x}_i(t) &= g_i(t) + \sum_j (-w_{ji}) l_j(t) \\ &= g_i(t) + \sum_j (-w_{ji}) \sum_k w_{jk} x_k(t), \end{aligned} \quad (28)$$

where \tilde{x}_i is the input of the i -th MC and g_i is the output of the corresponding glomerulus. Here, we assume the excitatory connections and the self-inhibitory feedback connections to have the same weights. The active MC is assumed to behave like the integrator described in (6). Both (27) and (28) can be written in matrix form, using the vectors \mathbf{g} , \mathbf{l} , $\tilde{\mathbf{x}}$ and \mathbf{x} which contain the output of the glomeruli, the activity levels of the GCs and the inputs and outputs of the MCs, respectively. Thus, the model of the GCs becomes $\mathbf{l}(t) = \mathbf{W} \mathbf{x}(t)$ and the model for the MCs with interconnected GCs follows as:

$$\tilde{\mathbf{x}}(t) = \mathbf{g}(t) - \mathbf{W}^T \mathbf{W} \mathbf{x}(t).$$

One potential task of the MOB is to decompose a mixture of chemical stimuli into its individual parts. Let \mathbf{o} be a vector of chemical stimuli of arbitrary dimension. The sensors of the olfactory system, the OSNs, react to the chemical components of this mixture and produce an output vector \mathbf{u} of dimension N . Here, we do not consider individual OSNs. Each entry of the vector can be thought of as the average output of the OSNs of one “type” since the purpose of the high

number of individual sensors is most likely noise reduction and redundancy. Thus, the dimension N corresponds directly to the number of glomeruli in the system. Say the olfactory system has an internal model of odor patterns and the problem is to find the most plausible mixture that causes the current (noisy) OSN output vector $\mathbf{u}(t)$. A smart way of solving this task would be to minimize the quadratic distance between $\mathbf{u}(t)$ and a vector created according to the internal model. Assuming linear mixing of the chemical stimuli, the “known” stimuli patterns can be collected as the columns of a matrix and we assume them to be reflected in the weight matrix \mathbf{W} . If the processing of the SACs (and PGCs) on the glomerulus level can be described by $\mathbf{g}(t) = \mathbf{W}^T \mathbf{u}(t)$, then we can see that the total system behaves just like the system defined by (7), explained above. Hence, under these assumptions if $\mathbf{u}(t)$ is almost constant for a long enough time to let the dynamic system converge and the synaptic weights contain the “known” or “learned” stimuli patterns, the system solves a NNLS optimization problem to find the most likely mixture of chemical stimuli encountered. More precisely, for constant $\mathbf{u}(t) = \mathbf{u}$ as $t \rightarrow \infty$, $\mathbf{x}(t)$ converges to the minimizer of the NNLS optimization problem (8), with the change of variable $\mathbf{A} = \mathbf{W}$ and $\mathbf{b} = \mathbf{u}$. Likewise, the system equations follow by a change of variable $\mathbf{A} = \mathbf{W}$ as (7).

We can make some observations stemming from the modeling. The weights in the model are non-negative, so that the weight matrix $\mathbf{W} \geq 0$, or equivalently $\mathbf{W} \in \mathbb{R}_{\geq 0}^{N \times N}$. Assuming \mathbf{W} is full rank, it is a basis for \mathbb{R}^N . It cannot, however, generate all points in $\mathbb{R}_{\geq 0}^N$ using only positive linear combinations, i.e., $\boldsymbol{\theta} = \mathbf{W}\boldsymbol{\gamma}$ with $\boldsymbol{\gamma} \in \mathbb{R}_{\geq 0}^N$. As said above, the vector \mathbf{u} is the noisy output of the OSNs: a noisy measurement. Say $\mathbf{u} = \mathbf{W}\mathbf{x}_0 + \boldsymbol{\eta}$, where \mathbf{x}_0 is the actual mixture and $\boldsymbol{\eta}$ is an additive noise vector. Then, although the solution \mathbf{x}^* of the NNLS problem is unique and $\boldsymbol{\eta}$ must be linearly dependent on \mathbf{W} , the optimal value $\frac{1}{2} \|\mathbf{W}\mathbf{x}^* - \mathbf{u}\|_2^2 \geq 0$. In contrast, the solution of a regular least-squares optimization problem without the non-negativity constraint would always yield an optimal value of zero in this case.

In [29], a similar modeling of the olfactory system is presented. However, it differs fundamentally from ours since there the GCs are modeled as active components whereas they are modeled statically here. Furthermore, a sparsity inducing cost function is assumed in [29], which is obsolete in our model due to the self-regularizing behavior induced by the non-negativity constraint.

B. Numerical Evaluation

To evaluate the capabilities of the dynamic olfactory system model we performed Monte Carlo simulations. Two *non-negative* data models were used in these simulations, one based on Gaussian random variables and the other based on rectangular random variables. Evidently, the performance of the dynamic olfactory system model must be evaluated for noisy data. In the following, the data models used in the Monte Carlo simulations are explained in detail. Note that we use the same data model in Section VI and so the data models are introduced for non-square matrices $\mathbf{W} \in \mathbb{R}^{M \times N}$, although $M = N$ in the olfactory system model.

Let \mathbf{x}_0 be the actual mixture, having exactly s nonzero entries at the indices contained in the sparse support set \mathcal{S} . Conversely, the indices where \mathbf{x}_0 is zero are in the complementary sparse support set, denoted as \mathcal{S}^c . The actual (noise free) input is $\mathbf{u}_0 = \mathbf{W}\mathbf{x}_0$. The dynamic system, however, is presented with a noisy input $\mathbf{u} = \mathbf{W}\mathbf{x}_0 + \boldsymbol{\eta}$, where $\boldsymbol{\eta}$ is an additive noise vector.

For the numerical evaluation, \mathbf{x}_0 , \mathbf{W} and $\boldsymbol{\eta}$ are drawn randomly. Here, the random matrix \mathbf{W} is entry-wise i.i.d. with expectation μ_w and variance σ_w^2 . The matrix \mathbf{W} is subsequently normalized, so that the euclidean column-norm is equal to one. Thus, each entry has expectation $\mathbb{E}[w_{ij}] = \frac{\mu_w}{\sqrt{M(\mu_w^2 + \sigma_w^2)}}$ and variance $\text{Var}[w_{ij}] = \frac{\sigma_w^2}{M(\mu_w^2 + \sigma_w^2)}$ after normalization. Each nonzero entry of the i.i.d. random vector \mathbf{x}_0 has expectation μ_x and variance σ_x^2 . Similarly, for the i.i.d. noise vector we have $\mathbb{E}[\eta_i] = 0$ and $\text{Var}[\eta_i] = \sigma_\eta^2$. Additionally, \mathbf{W} , \mathbf{x}_0 and $\boldsymbol{\eta}$ are entry-wise mutually independent.

We define the input signal-to-noise ratio (SNR) as $\frac{\mathbb{E}[\mathbf{u}_0^T \mathbf{u}_0]}{\mathbb{E}[\boldsymbol{\eta}^T \boldsymbol{\eta}]}$. Using the definitions from above, the input SNR can be given as (a detailed derivation can be found in the Appendix):

$$\frac{\mathbb{E}[\mathbf{u}_0^T \mathbf{u}_0]}{\mathbb{E}[\boldsymbol{\eta}^T \boldsymbol{\eta}]} = \frac{s}{M\sigma_\eta^2} \left(\sigma_x^2 + \frac{\mu_x^2 (s\mu_w^2 + \sigma_w^2)}{\mu_w^2 + \sigma_w^2} \right). \quad (29)$$

Irrespective of the data model, the nonzero entries in \mathbf{x}_0 are always drawn from a rectangular distribution with support $[0, \sqrt{12}]$, so that $\sigma_x^2 = 1$ and $\mu_x = \sqrt{3}$.

For the *rectangular* (rect) data model, the entries in \mathbf{W} and the entries in $\boldsymbol{\eta}$ are drawn from rectangular distributions. This corresponds to a situation where the support of input and noise values is naturally limited by the application, for example through saturation effects. The matrix entries are drawn with support $[0, \sqrt{12}]$, so that $\sigma_w^2 = 1$ and $\mu_w = \sqrt{3}$. Since the noise entries are zero mean, the support interval $[-\gamma, \gamma]$ can directly be determined for the chosen input SNR by solving (29) for σ_η^2 and observing that $\gamma = \sqrt{3}\sigma_\eta$.

In the *Gaussian* model the entries are drawn from Gaussian distributions for \mathbf{W} and $\boldsymbol{\eta}$. Since the matrix \mathbf{W} is assumed to be entry-wise non-negative, mean and variance are set to $\mu_w = 5$ and $\sigma_w^2 = 1$, respectively. Thus, the entries of \mathbf{W} are non-negative with very high probability. For $\boldsymbol{\eta}$, the parameter σ_η^2 is found analogously to the rectangular case.

For the numerical evaluation, 5000 Monte Carlo instances are generated. Note that in each instance the sparse support set \mathcal{S} , \mathbf{x}_0 , \mathbf{W} and $\boldsymbol{\eta}$ are drawn randomly as described above.

In Figure 4, on the left the mean relative error on the sparse support set \mathcal{S} is plotted for different input SNRs as well as different sparsities s for both non-negative data models. As can be seen, higher sparsities correspond to higher errors, which is to be expected. Also, the Gaussian data model generally performs worse than the rectangular one. Though relative errors of less than 10% can be reached, when the input SNR is high enough.

To show the denoising capabilities of this system, we evaluate the output SNR, which is defined as:

$$\mathbb{E} \left[\frac{[\mathbf{x}]_{\mathcal{S}}^T [\mathbf{x}]_{\mathcal{S}}}{[\mathbf{x}]_{\mathcal{S}^c}^T [\mathbf{x}]_{\mathcal{S}^c}} \right]. \quad (30)$$




Fig. 4. *left*: Mean relative error on the sparse support set \mathcal{S} for different input SNRs. *right*: Mean output SNR (see (30)) for different input SNRs.

There, $[\mathbf{x}]_{\mathcal{S}}$ is a vector comprised only of the entries of \mathbf{x} whose indices are contained in the set \mathcal{S} . Hence, the output SNR is the ratio between the power in the sparse support \mathcal{S} (desired signal) and the power in the complementary sparse support \mathcal{S}^c (noise) after recovery. The mean output SNR versus the input SNR is shown in Figure 4 on the right side. For both data models under low sparsities s , the output SNR is significantly higher than the input SNR. This suggests that the system possesses denoising capabilities.

Biological systems exhibit high fault tolerance and tend to be power efficient. Since the matrix \mathbf{W} models synaptic connections, we can test two effects by setting the smallest entries in the matrix to zero. On the one hand, we may simulate synaptic failures this way and thereby test the fault tolerance of the system. On the other hand, we may also test robustness and power efficient operation at the same time since maintaining a multitude of synaptic connections comes at a high cost for the body. In Figure 5, we show the mean output SNR evaluated for pruning ratios between 0 and 0.9. Here, a pruning ratio of 0.1 means that the smallest 10% of the matrix entries of \mathbf{W} are set to zero. Especially for the rectangular data model, we can see that even moderate pruning ratios ($\approx 0.5 - 0.7$) are acceptable since the output SNR stays higher than the input SNR.

In summary, these results indicate that a topologically simple neural network like the olfactory system could in principle tackle complicated signal recovery tasks, e.g., non-negative sparse approximation. It could greatly benefit from the inherent denoising capability and the fault tolerance operation of the presented network.

VI. APPLICATION: NON-NEGATIVE SPARSE APPROXIMATION

The second application we present is solving non-negative sparse approximation problems. Here, we discuss a more general model compared to the previous section including the very important class of underdetermined problems. We first introduce the model and present numerical evaluations afterwards.




Fig. 5. Mean output SNR for different pruning ratios of the matrix \mathbf{W} .

A. Model description

We consider non-negative sparse approximation problems as an application of the network presented in Section III. As introduced similarly in Section II-A, the task is to recover an s -sparse signal \mathbf{x}_0 from linear measurements. The non-zero entries of \mathbf{x}_0 are contained in the sparse support set \mathcal{S} and it is assumed that these entries are non-negative, i.e., $[\mathbf{x}_0]_{\mathcal{S}} \in \mathbb{R}_{\geq 0}^{|\mathcal{S}|}$, where $|\mathcal{S}|$ is the cardinality of the set \mathcal{S} . The measured signal \mathbf{b} is obtained from a noisy linear measurement process, which is described by $\mathbf{b} = \mathbf{A}\mathbf{x}_0 + \boldsymbol{\varepsilon}$ with $\mathbf{A} \in \mathbb{R}_{\geq 0}^{M \times N}$ and $\boldsymbol{\varepsilon} \in \mathbb{R}^M$. Typically, the number of measurements is assumed to be small so that $M \ll N$. Note that we assume for simplicity that all entries of the matrix \mathbf{A} are strictly greater than zero. This implies that the matrix fulfills the self-regularizing property [15] (see also Section II-A). Furthermore, we assume that the measured signal is element-wise non-negative, i.e., $\mathbf{b} \geq 0$, with very high probability. This holds if the noise $\boldsymbol{\varepsilon}$ is zero mean, the distribution of $\boldsymbol{\varepsilon}$ has finite support or tails with negligible contributions and the SNR is high enough.

To solve non-negative sparse approximation problems with the discontinuous dynamic system from Section III, the system is presented with the measurement at its input $\mathbf{u}(t) = \mathbf{b}$ and the system matrix equals the linear system model. The systems output $\mathbf{y}(t)$ will hold the recovered signal vector \mathbf{x} , once it has converged.

In the convergence analysis in Section IV, in some places it was required that the system matrix \mathbf{A} has full column-rank. Obviously, if $M < N$, as is typical in non-negative sparse approximation problems, the matrix cannot have full column-rank by definition. We discuss the consequences in the following. We see, that it does not affect the practical applicability of the network for the relevant problem class.

The solution of the NNLS problem is not unique if the matrix \mathbf{A} does not have full column-rank (compare Section IV-A). However, the KKT conditions of the discontinuous dynamic system and the NNLS problem are still equal. Moreover, since we assume that the matrix \mathbf{A} is entry-wise non-negative, the row-span of the matrix intersects the positive orthant and the self-regularizing property is fulfilled [15]. Thus, if no noise is present the solution is unique nevertheless and if noise is

present, the recovery error is bounded. So even if the solution of the NNLS problem would differ from the equilibrium point of the dynamic system in the noisy case, because the KKT conditions are equal, the error is bounded for both solutions.

For a given starting point $\mathbf{x}(0)$ and a given constant input $\mathbf{u}(t) = \mathbf{b}$, we can still split the system into linear subsystems (compare Section IV-B). The “active” subsystems corresponding to the index sets \mathcal{I}_j^+ are (locally) asymptotically stable since $-\mathbf{z}^T \mathbf{A}_j^+ \mathbf{A}_j^+ \mathbf{z} < 0$ for all $\mathbf{z} \neq \mathbf{0}$, due to the fact that $\mathbf{A}_j^+ > 0$. Hence, all state transitions remain unique and continuous. Also, the Caratheodory solution can be defined in the same way as before. The stability analysis can be performed using a given starting point $\mathbf{x}(0)$, a given constant input $\mathbf{u}(t) = \mathbf{b}$ and a corresponding equilibrium point \mathbf{x}_e (which depends on $\mathbf{x}(0)$ and \mathbf{b}). Since $\mathbf{A} > 0$, (23) still holds although $\mathbf{A}^T \mathbf{A}$ is only positive semi-definite. For a given tuple $(\mathbf{x}(0), \mathbf{u}(t) = \mathbf{b}, \mathbf{x}_e)$ and the condition on the solution the proof of the finite number of switches during convergence remains the same as well (c.f. Section IV-D).

Summarizing, even for the case that $M < N$ which prevents the matrix \mathbf{A} from having full column-rank, the system is (locally) asymptotically stable if $\mathbf{A} > 0$ and can thus be used to solve non-negative sparse approximation problems.

B. Numerical Evaluation

For the evaluation, we use the same non-negative data models introduced in Section V-B. Here, our main focus of the evaluation is to compare the recovery performance of the dynamic system (NNLS) with non-negative basis pursuit denoising (NNBPDN):

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \alpha \|\mathbf{x}\|_1 \\ & \text{subject to} && \mathbf{x} \geq \mathbf{0}. \end{aligned} \quad (31)$$

Note that in contrast to the NNLS problem, the NNBPDN problem has a regularization parameter α which must be determined in advance. For the evaluation, for each Monte Carlo instance, we solve the NNBPDN problem for 50 logarithmically spaced regularization parameters and pick the solution which has the smallest mean squared error (MSE) with respect to the actual signal vector \mathbf{x}_0 . This gives the NNBPDN problem an advantage since in practice one obviously cannot use this procedure and α must be fixed to some empirically determined value beforehand.

In Figure 6, the MSE of the recovered signal \mathbf{x} with respect to the actual signal vector \mathbf{x}_0 on the indices contained in \mathcal{S} , i.e., the MSE on the sparse support, is shown for different input SNRs. As can be seen, both methods differ only marginally with NNLS featuring a slightly lower error for a medium to high number of measurements M . Only for a low number of measurements and higher SNRs, NNBPDN has a smaller error on average.

Next, we evaluate the sparse support set recovery performance in Figure 7. Here, we check whether there exists a threshold with which the sparse support \mathcal{S} can be separated from the complementary sparse support \mathcal{S}^c , i.e., whether $\max([\mathbf{x}]_{\mathcal{S}^c}) < \min([\mathbf{x}]_{\mathcal{S}})$. We can see, that NNBPDN seems




Fig. 6. Comparison of the MSE with respect to the actual signal vector \mathbf{x}_0 on the sparse support set \mathcal{S} for different input SNRs, $N = 200$, $s = 5$ and using the uniform (rect) data model. Note the logarithmic scaling of the ordinate.

to be slightly superior as far as recovery of the support is concerned.




Fig. 7. Fraction of monte carlo instances in which clear recovery of the sparse support set \mathcal{S} is possible through thresholding for different input SNRs, $N = 200$, $s = 5$ and using the uniform (rect) data model.

To analyze the influence of the sparsity on the recovery performance, Figure 8 depicts the MSE on the support on the left and the support recovery fraction on the right both for different values of the sparsity s . As expected, the higher the sparsity the higher the error and hence also the lower the fraction of recoverable supports. Although NNBPDN performs slightly superior for most choices of the parameters, it is interesting to note that for low SNR situations the discontinuous network yields solutions that have a lower MSE on the support on average.

So far, all results were presented for the uniform (rect) data model. A comparison between the recovery performance of the uniform and the Gaussian data model is shown in Figure 9. Similar to the results in Section V-B, the Gaussian data model generally performs worse. The general effects discussed above are similarly present with the Gaussian data model, however.

Summarizing, we see that the recovery performance of




Fig. 8. Influence of the sparsity s for different input SNRs, $M = 50$, $N = 200$ for the uniform (rect) data model. *left*: MSE on sparse support \mathcal{S} . *right*: sparse support recovery by thresholding.




Fig. 9. Comparison of the two non-negative data models for an input SNR of 40 dB, $s = 5$, $M = 50$ and $N = 200$. *left*: MSE on sparse support \mathcal{S} . *right*: sparse support recovery by thresholding. Note the logarithmic scaling of the *left* ordinate.

the discontinuous network is very similar to the performance obtained by solving a NNBPDPN optimization problem. The NNBPDPN approach tends to recover the sparse support slightly better. For most choices of the parameters, the dynamic system features a marginally lower MSE on the support, this difference becomes more profound in low SNR situations. Nevertheless, we expect these differences to be smaller in practice since the regularization parameter α of the NNBPDPN was chosen as the best from a set of 50 values for each Monte Carlo instance, which cannot be done in practice. Thus, the discontinuous network shows a very promising recovery performance for non-negative sparse approximation applications.

C. Relation to existing work

As briefly summarized in Section I, several networks exist which are used in similar applications. Firstly, the most profound novelty of our presented system (7) is that the integrators are discontinuous and have a lower limit equal to

zero (or are individually upper- and lower-bounded, compare Section IV-E). In contrast to this, in related systems (cf. [6]–[8]), the output of the integrators is not constrained. Although for some of these systems, the output is directly processed by a nonlinear thresholding function before being connected to neighboring nodes, all of these systems still feature unbounded internal system states. This also influences the convergence behavior of said systems. The limited integrator in our system captures inherent limits, e.g., on neuronal firing rates, directly. However, careful stability analysis is required due to the discontinuous nature of the system. Secondly, our original application for which our neural network was investigated (see Section V-A) features self-inhibition that violates the necessary conditions on stability in classical continuous HNNs.

VII. CONCLUSION

In this work, we investigated a discontinuous neural network that is applied as a model of the mammalian olfactory system and more generally as a method for solving non-negative sparse approximation problems. The discontinuities in the system function of the neural network arise from the fact that the integrators of the system are bounded to having non-negative outputs. Hence, the integrators “switch” between being active if the output is greater than zero and being inactive if the output reaches zero. First of all, we showed that the presented neural networks features equilibrium points that correspond to the solutions of a general non-negative least squares problem. We then specified a Caratheodory solution to the system’s discontinuous differential equations and proceeded to prove that the network is stable, provided that the system matrix has full column-rank, using Lyapunov’s direct method. Moreover, we were able to show, under a mild condition on the equilibrium point, that the system performs only a finite amount of switches while converging. We then generalized the obtained results to the case where the integrators of the system have arbitrary lower- and upper limits. In this case, the neural network converges to the solution of a bounded least squares optimization problem. Two applications for the presented neural network were shown. Firstly, our original application was shown, where the neural network was derived from modeling of the mammalian olfactory system. We used this model to investigate a potential application of the olfactory system, which is the decomposition of a mixture of chemical stimuli from noisy neural responses. Based on our findings, we argued that the olfactory system may be capable of performing such a complex sparse recovery task. Secondly, we generalized the application to embrace non-negative sparse approximation problems. We showed that the stability of the neural networks is unaffected under this model, even if the system matrix does not possess full column-rank. Then, we compared the recovery performance of the network with a classical non-negative basis pursuit denoising algorithm. Since the recovery performance differs only marginally, the presented network can be successfully applied. Compared to the classical algorithm, the network is independent from the choice of a regularization parameter, which strongly affects the recovery performance of the classical non-negative basis pursuit denoising algorithm.

APPENDIX

INPUT SNR DERIVATION

Let \mathbf{I}_S denote a sparse identity matrix, i.e., a diagonal matrix with ones only at positions $\{(i, i) \mid i \in S\}$ and zeros otherwise. Likewise, let $\mathbf{1}_S$ denote a sparse one-vector containing ones solely at indices $\{i \mid i \in S\}$ and zeros otherwise. For the derivation of the input signal power $E[\mathbf{u}_0^T \mathbf{u}_0]$, we need $E[\mathbf{W}^T \mathbf{W}]$, $\text{Tr}[E[\mathbf{W}^T \mathbf{W}] \mathbf{I}_S]$ and $\mathbf{1}_S^T E[\mathbf{W}^T \mathbf{W}] \mathbf{1}_S$. Considering the independencies between the entries we get:

$$E[\mathbf{W}^T \mathbf{W}]_{ij} = E\left[\sum_{k=1}^M w_{ki} w_{kj}\right] = M E[w_{1i} w_{1j}]$$

$$\Rightarrow E[\mathbf{W}^T \mathbf{W}] = \frac{\mu_w^2}{\mu_w^2 + \sigma_w^2} \mathbf{1}\mathbf{1}^T + \frac{\sigma_w^2}{\mu_w^2 + \sigma_w^2} \mathbf{I}.$$

Based on this result it follows that

$$\text{Tr}[E[\mathbf{W}^T \mathbf{W}] \mathbf{I}_S] = s,$$

and

$$\begin{aligned} \mathbf{1}_S^T E[\mathbf{W}^T \mathbf{W}] \mathbf{1}_S &= \frac{\mu_w^2}{\mu_w^2 + \sigma_w^2} \mathbf{1}_S^T \mathbf{1}\mathbf{1}^T \mathbf{1}_S + \frac{\sigma_w^2}{\mu_w^2 + \sigma_w^2} \mathbf{1}_S^T \mathbf{1}_S \\ &= \frac{s^2 \mu_w^2 + s \sigma_w^2}{\mu_w^2 + \sigma_w^2}. \end{aligned}$$

Finally, we can calculate the input signal power $E[\mathbf{u}_0^T \mathbf{u}_0]$ using the results obtained above:

$$\begin{aligned} E[\mathbf{u}_0^T \mathbf{u}_0] &= E[\mathbf{x}_0^T \mathbf{W}^T \mathbf{W} \mathbf{x}_0] = E[\text{Tr}[\mathbf{W}^T \mathbf{W} \mathbf{x}_0 \mathbf{x}_0^T]] \\ &= \text{Tr}[E[\mathbf{W}^T \mathbf{W}] E[\mathbf{x}_0 \mathbf{x}_0^T]] \\ &= \text{Tr}[E[\mathbf{W}^T \mathbf{W}] (\text{Cov}[\mathbf{x}_0] + E[\mathbf{x}_0] E[\mathbf{x}_0]^T)] \\ &= \text{Tr}[E[\mathbf{W}^T \mathbf{W}] (\sigma_x^2 \mathbf{I}_S + \mu_x^2 \mathbf{1}_S \mathbf{1}_S^T)] \\ &= \sigma_x^2 \text{Tr}[E[\mathbf{W}^T \mathbf{W}] \mathbf{I}_S] + \mu_x^2 \mathbf{1}_S^T E[\mathbf{W}^T \mathbf{W}] \mathbf{1}_S \\ &= \sigma_x^2 s + \frac{\mu_x^2 (s^2 \mu_w^2 + s \sigma_w^2)}{\mu_w^2 + \sigma_w^2}. \end{aligned}$$

The noise power is simply $E[\boldsymbol{\eta}^T \boldsymbol{\eta}] = M \sigma_\eta^2$. Thus, the input SNR follows as (29).

REFERENCES

- [1] D. L. Donoho, "Compressed sensing," *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [2] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [3] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the National Academy of Sciences*, vol. 79, no. 8, pp. 2554–2558, 1982.
- [4] J. J. Hopfield and D. W. Tank, "'Neural' computation of decisions in optimization problems," *Biological Cybernetics*, vol. 52, no. 3, pp. 141–152, 1985.
- [5] U.-P. Wen, K.-M. Lan, and H.-S. Shih, "A review of Hopfield neural networks for solving mathematical programming problems," *European Journal of Operational Research*, vol. 198, no. 3, pp. 675–687, 2009.
- [6] A. Bouzerdoum and T. Pattison, "Neural network for quadratic optimization with bound constraints," *IEEE Transactions on Neural Networks*, vol. 4, no. 2, pp. 293–304, 1993.
- [7] Y. Xia, "A new neural network for solving linear and quadratic programming problems," *IEEE Transactions on Neural Networks*, vol. 7, no. 6, pp. 1544–1548, 1996.
- [8] C. J. Rozell, D. H. Johnson, R. G. Baraniuk, and B. A. Olshausen, "Sparse Coding via Thresholding and Local Competition in Neural Circuits," *Neural Computation*, vol. 20, no. 10, pp. 2526–2563, 2008.
- [9] A. Balavoine, J. Romberg, and C. Rozell, "Convergence and Rate Analysis of Neural Networks for Sparse Approximation," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 23, no. 9, pp. 1377–1389, 2012.
- [10] E. J. Candes, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [11] E. Candes and M. Wakin, "An Introduction To Compressive Sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, 2008.
- [12] D. Chen and R. J. Plemmons, "Nonnegativity constraints in numerical analysis," in *Symposium on the Birth of Numerical Analysis*, 2009, pp. 109–140.
- [13] A. Bruckstein, M. Elad, and M. Zibulevsky, "On the Uniqueness of Non-negative Sparse Solutions to Underdetermined Systems of Equations," *IEEE Transactions on Information Theory*, vol. 54, no. 11, pp. 4813–4820, 2008.
- [14] M. Wang and A. Tang, "Conditions for a unique non-negative solution to an underdetermined system," in *47th Annual Allerton Conference on Communication, Control, and Computing*, 2009, pp. 301–307.
- [15] M. Slawski and M. Hein, "Non-negative least squares for high-dimensional linear models: Consistency and sparse recovery without regularization," *Electronic Journal of Statistics*, vol. 7, pp. 3004–3056, 2013.
- [16] J. Adamy, "Grundlagen nichtlinearer Systeme," in *Nichtlineare Regelungen*, ser. Springer-Lehrbuch. Berlin Heidelberg, DE: Springer, 2009, pp. 1–43.
- [17] J. J. E. Slotine and W. Li, *Applied nonlinear control*. Englewood Cliffs, NJ, US: Prentice Hall, 1991.
- [18] H. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River, NJ, US: Prentice Hall, 2002.
- [19] S. Sastry, *Nonlinear Systems: Analysis, Stability and Control*, ser. Interdisciplinary applied mathematics. New York, NY, US: Springer, 1999, vol. 10.
- [20] J. Cortes, "Discontinuous dynamical systems," *Control Systems, IEEE*, vol. 28, no. 3, pp. 36–73, 2008.
- [21] S. Boyd and L. Vandenberghe, *Convex Optimization*, 7th ed. New York, NY, US: Cambridge University Press, 2004.
- [22] A. Björck, "Constrained Least Squares Problems," in *Numerical Methods for Least Squares Problems*, ser. Other Titles in Applied Mathematics. Philadelphia, PA, US: Society for Industrial and Applied Mathematics, 1996, vol. 51, pp. 187–213.
- [23] J. Lunze, "Beschreibung und Verhalten von Mehrgrößensystemen," in *Regelungstechnik 2*, 7th ed., ser. Springer-Lehrbuch. Berlin Heidelberg, DE: Springer, 2013, vol. 2, pp. 15–62.
- [24] L. B. Buck, "Information Coding in the Vertebrate Olfactory System," *Annual Review of Neuroscience*, vol. 19, no. 1, pp. 517–544, 1996.
- [25] G. Laurent, "A Systems Perspective on Early Olfactory Coding," *Science*, vol. 286, no. 5440, pp. 723–728, 1999.
- [26] P. Mombaerts, "Axonal Wiring in the Mouse Olfactory System," *Annual Review of Cell and Developmental Biology*, vol. 22, no. 1, pp. 713–737, 2006.
- [27] R. I. Wilson and Z. F. Mainen, "Early events in olfactory processing," *Annual Review of Neuroscience*, vol. 29, no. 1, pp. 163–201, 2006.
- [28] V. N. Murthy, "Olfactory Maps in the Brain," *Annual Review of Neuroscience*, vol. 34, no. 1, pp. 233–258, 2011.
- [29] A. Koulakov and D. Rinberg, "Sparse Incomplete Representations: A Potential Role of Olfactory Granule Cells," *Neuron*, vol. 72, no. 1, pp. 124–136, 2011.